A COMPARATIVE STUDY OF ALTERNATIVE EXTREME-VALUE VOLATILITY ESTIMATORS

TURAN G. BALI*
DAVID WEINBAUM

Recent advances in econometric methodology and newly available sources of data are used to examine empirically the performance of the various extreme-value volatility estimators that have been proposed over the past two decades. Overwhelming support is found for the use of extreme-value estimators when computing daily volatility measures across all assets: Daily extreme-value volatility estimators are both less biased and substantially more efficient than the traditional close-to-close estimator. In the case of weekly and monthly measures, the results still suggest that extreme-value estimators are appropriate, but the evidence is more mixed. © 2005 Wiley Periodicals, Inc. Jrl Fut Mark 25:873–892, 2005

The authors thank Menachem Brenner, Peter Christoffersen, Robert Engle, Stephen Figlewski, Peter Frost, Robert Jarrow, Robert Webb (the Editor), and an anonymous referee for helpful comments and discussions. We acknowledge Kai Li's contributions to an earlier version of this article. All errors are solely the authors' responsibility.

*Correspondence author, Zicklin School of Business, Baruch College, City University of New York, 17 Lexington Avenue, Box E-621, New York, New York 10010; e-mail: Turan_Bali@baruch.cuny.edu

Received July 2004; Accepted December 2004

■ *Turan G. Bali is an Associate Professor of Economics and Finance at Baruch College at the Zicklin School of Business at the City University of New York in New York.*

■ *David Weinbaum is an Assistant Professor of Finance at the Johnson Graduate School of Management at Cornell University in Ithaca, New York.*

The Journal of Futures Markets, Vol. 25, No. 9, 873–892 (2005)
Published online in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/fut.20169

INTRODUCTION

Asset return volatility is important in the theory and practice of financial economics. Accurate measures and good forecasts of volatility are therefore critical. This article compares the performance of various methods of estimating volatility from daily data (opening, closing, high, and low prices). The analysis focuses on measuring volatility as opposed to forecasting it, so the work is relevant for applications in which an efficient volatility estimate (e.g., event studies) rather than an efficient volatility forecast (e.g., risk management) is called for.[1]

The traditional volatility estimator is the sample standard deviation of close-to-close returns computed over the relevant horizon. An alternative approach is to use the information contained in the highest and lowest prices observed during the trading day. A number of these so-called extreme-value volatility estimators have been proposed over the past two decades, their main advantage over the traditional estimator being that they are, at least in theory, significantly more efficient: It can be shown that their sampling variance is 5 to 14 times lower than that of the traditional estimator.[2]

Some caution must be exercised when employing these estimators. First, all extreme-value volatility estimators are derived under the assumption that the price of the asset follows geometric Brownian motion. Violations of this assumption will not only affect the efficiency gains that these estimators afford in theory, but may also cause them to be biased; the direction of the bias will depend on the empirical distribution of the price of the asset. Second, microstructure effects may also contribute to a bias in these estimators. Because prices are not observed continuously, the reported high and low prices may not be the true high and low prices. Thus discrete trading may impart a downward bias in these estimators. If transaction prices (as opposed to quotes) are used, then there is an offsetting tendency for a positive bias in extreme-value volatility estimators due to the bid–ask spread: the daily high is most likely to occur at the ask, and the low will usually be at the bid.

These observations have spurred interest in testing the performance of extreme-value volatility estimators. The major difficulty that the

[1]For a recent application of extreme-value volatility estimators in the context of an event study, see, for example, Brown and Hartzell (2001).

[2]The various extreme-value volatility estimators are presented in the next section. There are also other extreme-value approaches to estimating volatility. For example, the EVT approach introduced by Bali (2003) and Bali and Neftci (2003) is based on the maximum-likelihood parameter estimation of the generalized Pareto and generalized extreme-value distributions.

literature has had to deal with in this respect is that volatility is not directly observable. Researchers have addressed this in two ways. One strand of the literature uses simulations and theoretical arguments. For example, Marsh and Rosenfeld (1986) and Cho and Frees (1988) propose theoretical models to analyze the impact of discrete trading on the estimation of volatility. Marsh and Rosenfeld (1986) find that discrete trading does not bias the traditional estimator (although it reduces its efficiency), and that it biases extreme-value estimators downward and causes them to become less efficient. Cho and Frees (1988) find that discrete trading causes the traditional estimator to be upward biased. Garman and Klass (1980) recognize that their extreme-value estimator is biased downward in the presence of discrete trading, and show that the traditional estimator is upward biased in this context, but only slightly so. In simulated data, Rogers, Satchell, and Yoon (1994) find that the question of which estimator to use depends on the distributional assumption that one makes about returns. A second strand of the literature uses actual daily securities data and relies on additional assumptions to get around the problem that, in such data, volatility is not precisely measurable. In this vein, Beckers (1983) and Wiggins (1991) study the performance of the Parkinson (1980) extreme-value volatility estimator using the traditional estimator as a benchmark. Their overall conclusion is that extreme-value estimators are "better" in a sense elaborated upon in these articles.

This work differs from previous research in this area along three dimensions. First, a new approach is used to address the fact that volatility is not observable: Recent advances in econometric methodology and richer sources of data are used. In a recent series of articles, Andersen and Bollerslev (1998), Andersen, Bollerslev, Diebold, and Labys (2001, in press), Andreou and Ghysels (2002), and Barndorff-Nielsen and Shephard (2002) suggest the sum of squared high-frequency intraday returns as a precise measure of volatility. Following these researchers, high-frequency data are used on very liquid and actively traded assets to construct realized-volatility measures. Volatility is thus treated as observed (rather than latent), and the various estimators are compared directly, with actual securities data and without making distributional assumptions about asset returns. Second, this analysis recognizes explicitly that the traditional close-to-close estimator may be biased (even in the absence of microstructure effects). This is in contrast to Beckers (1983) and Wiggins (1991), who assume that the traditional estimator is unbiased and estimate the empirical bias in Parkinson's extreme-value estimator relative to the traditional estimator. Third, a horse race is run among all the extreme-value estimators that have appeared in the past two decades, whereas previous research has only studied a subset of these.

One might wonder why one would be at all interested in extreme-value volatility estimators when high-frequency realized volatility estimators are available. This is because there are important situations in which intraday data are unavailable or unreliable (e.g., because of thin trading). Although accurate measures of realized volatility require high frequency, 5-minute returns, Engle and Patton (in press) use TAQ data to group NYSE stocks into trading frequency deciles. They find that the median intertrade time in the first decile is over 90 minutes, in the second decile it is 27 minutes, and in the fourth decile it is 9 minutes. The median intertrade time does not drop below 5 minutes until the sixth decile. This thin trading problem is avoided here altogether by focusing on S&P 500 futures contracts that are very actively traded, and the analysis is augmented with currency data as a robustness check.

The rest of this article is organized as follows. The next section describes the various extreme-value volatility estimators and shows why it is generally erroneous to think of the traditional estimator as unbiased. It also discusses the construction of the realized-volatility measures, which are central to this work. The data and the empirical results are presented next, followed by a conclusion.

VOLATILITY ESTIMATION

The various volatility estimators that have appeared in the literature are now introduced, and some of their properties discussed.

Extreme-Value Volatility Estimators

All extreme-value volatility estimators are derived under the assumption that S_t, the price of the asset at time t, follows geometric Brownian motion, that is, that it satisfies the following stochastic differential equation:

$$dS_t = \mu S_t \, dt + \sigma S_t \, dW_t \tag{1}$$

where μ and $\sigma > 0$ are constants, and W_t is a standard Brownian motion. The following notation will be needed. Let O_t, C_t, H_t, and L_t

denote, respectively, the opening, closing, high, and low prices on day t, and n the number of days in the sample. To maintain consistency with the theoretical literature on extreme-value estimators, the presentation that follows is in terms of variances, not volatilities. In the empirical section, however, annualized volatilities are used, as the scale of these is more readily interpretable.

The traditional, or close-to-close, estimator of variance for a driftless[3] security is given by

$$\hat{\sigma}^2_{\mathrm{cc}} = \frac{1}{n}\sum_{t=1}^{n}\left(\ln\frac{C_t}{C_{t-1}}\right)^2, \qquad n \geq 1 \tag{2}$$

A mean-adjusted variant of this estimator is given by the sample standard deviation,

$$\hat{\sigma}^2_{\mathrm{acc}} = \frac{1}{n-1}\sum_{t=1}^{n}\left(\ln\frac{C_t}{C_{t-1}}\right)^2 - \frac{[\ln(C_n/C_0)]^2}{n(n-1)}, \qquad n > 1 \tag{3}$$

Parkinson (1980) introduces the following extreme-value estimator for a driftless security:

$$\hat{\sigma}^2_p = \frac{1}{4n\ln 2}\sum_{t=1}^{n}\left(\ln\frac{H_t}{L_t}\right)^2, \qquad n \geq 1 \tag{4}$$

The efficiency of an (unbiased) estimator $\hat{\sigma}^2$ relative to the traditional estimator $\hat{\sigma}^2_{\mathrm{cc}}$ is defined in the usual manner by the ratio $Var(\hat{\sigma}^2_{\mathrm{cc}})/Var(\hat{\sigma}^2)$. It can be shown that if the stock price follows Equation (1) with $\mu = \sigma^2/2$, and if trading is continuous and continuously monitored, then the Parkinson (1980) estimator $\hat{\sigma}^2_p$ is about five times more efficient than the traditional estimator (i.e., its sampling variance is about five times lower).

Under the assumptions of Parkinson (1980), Garman and Klass (1980) construct a minimum variance unbiased estimator that simultaneously uses the opening, closing, high, and low prices:

$$\hat{\sigma}^2_{\mathrm{GK}} = \frac{1}{n}\sum_{t=1}^{n}\left(0.511\left(\ln\frac{H_t}{L_t}\right)^2 - 0.019\left[\ln\left(\frac{C_t}{O_t}\right)\ln\left(\frac{H_tL_t}{O_t^2}\right) - 2\ln\left(\frac{H_t}{O_t}\right)\ln\left(\frac{L_t}{O_t}\right)\right] - 0.383\left[\ln\frac{C_t}{O_t}\right]^2\right), \qquad n \geq 1 \tag{5}$$

[3]By driftless, it is meant that the logarithmic price process is driftless; that is, $\mu = \sigma^2/2$. This is because the solution to (1) is given by $S_t = S_0 \exp\{(\mu - \frac{1}{2}\sigma^2)t + \sigma W_t\}$.

This estimator is theoretically 7.4 times more efficient than the traditional estimator, but still maintains the assumption that $\mu = \sigma^2/2$. Rogers and Satchell (1991) relax this assumption and propose the following estimator:

$$\hat{\sigma}^2_{\mathrm{RS}} = \frac{1}{n}\sum_{t=1}^{n}\left(\ln\left(\frac{H_t}{C_t}\right)\ln\left(\frac{H_t}{O_t}\right) + \ln\left(\frac{L_t}{C_t}\right)\ln\left(\frac{L_t}{O_t}\right)\right), \qquad n \geq 1 \quad (6)$$

which has the desirable property that it is independent of the drift μ. Because $\hat{\sigma}^2_{\mathrm{RS}}$ is a member of the unbiased quadratic class of Garman and Klass (1980), $\hat{\sigma}^2_{\mathrm{GK}}$ will outperform it if $\mu = \sigma^2/2$, but will go astray if this is not the case. Rogers and Satchell (1991) also propose an adjustment that is designed to take into account the fact that one may not be able to continuously monitor the stock price. Their adjusted estimator is the positive root of the following quadratic equation:

$$\hat{\sigma}^2_{\mathrm{ARS}} = \frac{0.5594}{N^{\mathrm{obs}}}\hat{\sigma}^2_{\mathrm{ARS}} + \frac{0.9072}{\sqrt{N^{\mathrm{obs}}}}\ln\left(\frac{H_t}{L_t}\right)\hat{\sigma}_{\mathrm{ARS}} + \hat{\sigma}^2_{\mathrm{RS}} \quad (7)$$

where N^{obs} denotes the number of observations of the price during the trading day.[4] They propose a similar adjustment to the Garman and Klass (1980) estimator; the adjusted estimator is the positive root of the following equation:

$$\begin{aligned}\hat{\sigma}^2_{\mathrm{AGK}} = {} & 0.511\left[\left(\ln\frac{H_t}{L_t}\right)^2 + \frac{0.9709}{N^{\mathrm{obs}}}\hat{\sigma}^2_{\mathrm{AGK}} + \frac{1.8144}{\sqrt{N^{\mathrm{obs}}}}\left(\ln\frac{H_t}{L_t}\right)\hat{\sigma}_{\mathrm{AGK}}\right] \\ & + 0.038\left[\ln\left(\frac{H_t}{O_t}\right)\ln\left(\frac{L_t}{O_t}\right) - \frac{0.2058}{N^{\mathrm{obs}}}\hat{\sigma}^2_{\mathrm{AGK}} - \frac{0.4536}{\sqrt{N^{\mathrm{obs}}}}\left(\ln\frac{H_t}{L_t}\right)\hat{\sigma}_{\mathrm{AGK}}\right] \\ & - 0.019\ln\left(\frac{C_t}{O_t}\right)\ln\left(\frac{H_tL_t}{O_t^2}\right) - 0.383\left(\ln\frac{C_t}{O_t}\right)^2 \end{aligned} \quad (8)$$

Finally, Yang and Zhang (2000) propose a minimum-variance unbiased estimator that is independent of the drift μ of the asset price process.[5]

[4]For a variance estimator over more than 1 day, simply take the arithmetic average of the $\hat{\sigma}^2_{\mathrm{ARS}}$'s over the interval of interest.

[5]This analysis ignores two extreme-value volatility estimators. Ball and Torous (1984) derive a maximum-likelihood estimator (which is asymptotically efficient) for the case $\mu = \sigma^2/2$, but their estimator has no closed-form solution. Kunitomo (1992) derives an extreme-value estimator that is about twice as efficient as Parkinson's estimator, but the Kunitomo estimator is based on the range of a Brownian bridge constructed from the price process, which implies that it cannot be computed from daily data.

Their practical estimator is given by

$$\hat{\sigma}^2_{\mathrm{YZ}} = \frac{1}{n-1}\sum_{t=1}^{n}\left(\ln\left(\frac{O_t}{C_{t-1}}\right) - \bar{o}\right)^2 + k\frac{1}{n-1}\sum_{t=1}^{n}\left(\ln\left(\frac{C_t}{O_t}\right) - \bar{c}\right)^2 + (1-k)\hat{\sigma}^2_{\mathrm{RS}}, \qquad n > 1 \tag{9}$$

where $\bar{o} = (1/n)\sum_{t=1}^{n}\ln(O_t/C_{t-1})$, $\bar{c} = (1/n)\sum_{t=1}^{n}\ln(C_t/O_t)$, and $k = 0.34/[1.34 + (n+1)/(n-1)]$.

Except for the Yang and Zhang (2000) estimator, $\hat{\sigma}^2_{\mathrm{YZ}}$, extreme-value estimators do not incorporate an estimate of overnight (i.e., closed market) variance: They ignore the fact that O_t, the opening price on day t, is in general different from C_{t-1}, the previous closing price. Inspection of Equation (9) reveals that $\hat{\sigma}^2_{\mathrm{YZ}}$ is simply the sum of the estimated overnight variance [the first term on the right-hand side of Equation (9)] and the estimated open market variance (which is a weighted average of the open-market return sample variance and the Rogers & Satchell, 1991, drift-independent estimator, where the weights are chosen so as to minimize the variance of the estimator). The resulting estimator therefore explicitly incorporates a term for closed-market variance; similar adjustments could be made to the other variance estimators; see Garman and Klass (1980) for details. In this article overnight variances are not estimated, because obviously high-frequency data do not allow the measurement of closed-market realized variance. Most of the empirical work centers on the foreign exchange market; this market is open 24 hours per day and is therefore unaffected by this problem. With equity data, open-market analogues of (2), (3) and (9) are used, for example, open-market sample variance

$$\hat{\sigma}^2_{\mathrm{aoc}} = \frac{1}{n-1}\sum_{t=1}^{n}\left[\ln\left(\frac{C_t}{O_t}\right) - \bar{c}\right]^2 \tag{10}$$

is used in lieu of (3). This is not expected to have a significant impact on the results: Any additional term for overnight variance would be the same for the various extreme-value estimators and the benchmark derived from high-frequency data, leaving the difference between these two quantities unaffected.[6]

[6]Strictly speaking, this is only true for variances. Because the empirical work centers on volatilities, a small bias would arise because of a Jensen term.

The Biasedness of the Traditional Estimator

Although it is well known that extreme-value estimators are derived under strong distributional assumptions and may thus be biased, the assumption is often made in the literature that the traditional estimator provides unbiased estimates of volatility regardless of the data-generating process. That this is not true in general, quite aside from any microstructure effects, is the focus of this section. For ease of exposition, assume that the asset price process is a diffusion,

$$dS_t = \mu_t S_t \, dt + \sigma_t S_t \, dW_t \tag{11}$$

where μ_t and σ_t are adapted and bounded processes. Then the logarithmic asset price process is given by

$$d \ln S_t = (\mu_t - \tfrac{1}{2}\sigma_t^2) \, dt + \sigma_t \, dW_t \tag{12}$$

It then follows that $E[(d \ln S_t)^2] = Var[d \ln S_t] = \sigma_t^2 \, dt$, so squared infinitesimal returns provide unbiased variance estimates. Consider now the τ-period continuously compounded return $r_t(\tau) = \ln(S_{t+\tau}) - \ln(S_t)$:

$$r_t(\tau) = \int_t^{t+\tau} (\mu_s - \tfrac{1}{2}\sigma_s^2) \, ds + \int_t^{t+\tau} \sigma_s \, dW_s \tag{13}$$

Standard arguments yield $\int_t^{t+\tau} \sigma_s^2 \, ds$ as the quadratic variation. This expression defines the so-called integrated volatility, which is central to option pricing theory under stochastic volatility, and is the subject of much recent empirical and theoretical work, for example, Andersen, Bollerslev, Diebold, and Ebens (2001), Andersen, Bollerslev, Diebold, and Labys (2001, in press), Andreou and Ghysels (2002), and Barndorff-Nielsen and Shephard (2002). Clearly, in general, $Var[r_t(\tau)] \neq \int_t^{t+\tau} \sigma_s^2 \, ds$. The equality would obtain, for example, if $\int_t^{t+\tau} (\mu_s - \frac{1}{2}\sigma_s^2) \, ds$ were nonstochastic. In general, the relation between $Var[r_t(\tau)]$ and $\int_t^{t+\tau} \sigma_s^2 \, ds$ depends upon the data-generating process. Consider a specific example in which volatility is actually constant: Assume that returns are predictable and log prices follow the trending Ornstein-Uhlenbeck process:

$$d \ln S_t = (-\gamma(\ln S_t - \mu t) + \mu) \, dt + \sigma \, dW_t \tag{14}$$

where γ, μ, and σ are positive constants. Log prices are the sum of a zero-mean stationary autoregressive Gaussian process and a deterministic linear trend. This process has been studied by Lo and Wang (1995) in the context of option pricing. It can be shown that

$$Var[r_t(\tau)] = \frac{\sigma^2}{\gamma}[1 - e^{-\gamma\tau}] \tag{15}$$

or, equivalently,

$$\sigma^2 = \frac{Var[r_t(\tau)]}{\tau} \frac{\gamma\tau}{1 - e^{-\gamma\tau}} \tag{16}$$

which implies that under the trending Ornstein-Uhlenbeck specification, the sample variance of continuously compounded returns is not an unbiased estimator of σ^2. Note, however, that the adjustment factor on the right-hand side of (16) vanishes in the continuous-time limit:

$$\lim_{\tau\to 0} \frac{\gamma\tau}{1 - e^{-\gamma\tau}} = \lim_{\tau\to 0} \frac{\gamma}{\gamma e^{-\gamma\tau}} = 1 \tag{17}$$

which is consistent with the heuristic discussion of the unbiasedness of the sample variance estimator for infinitesimal returns. In other words, sample variance is unbiased in the limit, as sampling frequency increases without bound; however, in general, sample variance need not be an unbiased estimator of (instantaneous) variance σ^2.

Measuring Realized Volatility

This approach to comparing the performance of the various volatility estimators requires that precise volatility measures be constructed. To this end, high-frequency data on very liquid and actively traded assets are used to construct measures of realized volatility as in, for example, Andersen, Bollerslev, Diebold, and Ebens (2001), Andersen, Bollerslev, Diebold, and Labys (2001, in press), and Barndorff-Nielsen and Shephard (2002). They show that under the assumption that the logarithmic asset price process is a special semimartingale, the sum of squares of discretely sampled continuously compounded returns computed from equally spaced observations converges uniformly in probability to the quadratic variation of the process (which is interpretable as realized cumulative instantaneous variability), as the sampling frequency increases without bound. The generality of the special semimartingale assumption is worth emphasizing; it assumes little more than absence of arbitrage opportunities. It does not require that the process be Markov, nor does it rule out jumps; for example, it encompasses Merton's (1976) jump diffusion model as well as pure diffusion models that are central to much stochastic volatility option pricing theory. For example, if the asset price process is a diffusion as in (11), then the sum of squares of discretely sampled returns over the time interval $[t_1, t_2]$ converges in probability to $\int_{t_1}^{t_2} \sigma_s^2\, ds$ as the sampling frequency increases without bound. In short, with high-frequency data, it is theoretically possible to construct volatility measures that are model- and error-free, the argument being essentially that observation of the full (i.e., continuous) sample path allows one to infer the true volatility process from the quadratic variation. Of course, in practice, the full sample path is not observable (one only has a discrete sequence of prices recorded at closely spaced times); nor is it desirable to sample the data at the highest available frequency, which would introduce a number of microstructure biases. As is standard in this literature, a 5-minute sampling interval is used as an appropriate trade-off between microstructure effects and measurement error.

EMPIRICAL RESULTS

Data

Four sets of high-frequency data are used: a stock market index futures contract and three currencies. All are very liquid and actively traded instruments. The S&P 500 index futures series consists of time-stamped transactions data from the CME;[7] the sample period is January 1989–December 2003. The exchange rates considered are the deutsche mark/U.S. dollar, Japanese yen/U.S. dollar, and U.S. dollar/British pound spot exchange rates. These are among the most actively traded and quoted currencies. The data were obtained from Olsen and Associates, Zurich, and cover the period December 1986–August 2003. The data are derived from all the bid–ask interbank quotes (not transaction prices) that have appeared on the Reuters screen over the sample period. Prices are obtained by averaging the log bid and log ask. (See Andersen, Bollerslev, Diebold, & Labys, 2001, for a complete description of the data set.)

[7]These are prices recorded by exchange personnel who observe the pit and post the most recent transaction price. The observers record every change in price, but not successive trades at the same price.

TABLE I
Descriptive Statistics for Realized Volatility

	Mean	*Standard deviation*	*Skewness*	*Kurtosis*
		Panel A: Equity index futures		
S&P 500 Futures	11.8964	5.8494	2.7163	13.8731
		Panel B: Exchange rates		
DM/US	10.7954	3.9798	1.7947	5.5134
JY/US	11.6321	4.8653	3.3108	32.0812
BP/US	9.9254	3.4572	1.8966	7.0987

Note. Entries are sample moments computed for daily realized volatility for the S&P 500 index futures contract and spot exchange rates on the Deutsche mark, the Japanese yen, and the British pound. The sample periods are from January 3, 1989 to December 30, 2003 for the S&P 500 futures contract, and from December 3, 1986 to August 30, 2003 for the exchange rates. Skewness is the estimate of the skewness measure $\gamma_1 = \mu_3/\mu_2^{3/2}$, and kurtosis is the estimate of the kurtosis measure $\gamma_2 = \mu_4/\mu_2^2 - 3$. Both are zero for normal random variables.

Table I provides summary statistics for the unconditional distribution of the realized daily volatility series computed for these data. The mean volatility of the S&P 500 index futures contract is 11.90%. Average volatilities for the currencies range from 9.93% for the British pound to 11.63 percent for the Japanese yen. The standard deviations reported in the second column show that realized volatility exhibits significant time series variation. The third and fourth columns indicate that realized volatility is extremely right-skewed and leptokurtic, which is consistent with previous research.

Comparison Criteria

In order to analyze the empirical performance of the various estimators presented in the previous section, one must define finite sample criteria upon which meaningful comparisons may be based. Although unbiasedness is a desirable attribute, it is rarely used by itself as an estimation criterion. Estimators are often compared on the basis of their mean-squared error. Let σ_t denote the volatility realized during period t, the mean-squared error of an estimator $\hat{\sigma}_t$ is

$$\begin{aligned} \mathrm{MSE}(\hat{\sigma}_t) &= E[(\hat{\sigma}_t - \sigma_t)^2] \\ &= (E[\hat{\sigma}_t - \sigma_t])^2 + Var[\hat{\sigma}_t - \sigma_t] \end{aligned} \tag{18}$$

thus the mean-squared error equals the square of the bias (mean difference between the estimator and the parameter) plus the variance of the difference (between the estimator and the parameter). In a decision-theoretic setting, minimizing the mean-square error is equivalent to minimizing the expected loss associated with a loss function that is proportional to the square of the difference between the estimate and the true parameter. It should be noted that the quantities in Equation (18) are based on squares of variance estimators, so the fourth moments of the data are involved. To check that the presence of outliers is not driving the results, the mean absolute difference between the estimator and the true parameter is also considered:

$$\mathrm{MAD}(\hat{\sigma}_t) = E[|\hat{\sigma}_t - \sigma_t|] \tag{19}$$

Finally, to give some sense of the magnitude of the bias associated with a given volatility estimator, the relative bias is also reported:

$$\text{Prop. Bias} = E\left(\frac{\hat{\sigma}_t - \sigma_t}{\sigma_t}\right) \tag{20}$$

Results

Table II reports sample estimates of the comparison criteria defined above, along with standard errors in parentheses. These and other statistics (such as estimates of the covariance matrix of the criteria) were obtained with the use of standard bootstrap techniques; for example, see Efron and Tibshirani (1986).

Panel A of Table II gives the results for the S&P 500 index futures contract, Panels B, C, and D are for the Deutsche mark/dollar exchange rate, the yen/dollar exchange rate, and the dollar/pound exchange rate, respectively. Volatility estimators are computed over nonoverlapping estimation windows of 1 day, 1 week (5 trading days), and 1 month (21 trading days). It was found to be useful to segment the results into daily volatility estimation, and weekly and monthly volatility estimation.

Daily Volatility Estimation

Several important results emerge from the analysis. First, the traditional estimator is significantly biased in all four data sets. In other words, consistent with the discussion in the previous section, it was found that squared returns do not provide unbiased estimates of the *ex post* realized volatility. Of particular interest, across the four data sets, extreme-value

TABLE II
Estimated Comparison Criteria

		Prop. Bias	*Bias*	*Variance*	*MSE*	*MAD*
		Panel A. S&P 500 futures				
$\hat{\sigma}_{oc}$	1 day	−0.2084	−2.2402	67.6112	72.6297	6.3945
		(0.0105)	(0.1405)	(4.0632)	(3.8411)	(0.1002)
	5 days	−0.0355	−0.3169	19.8905	19.9910	3.0978
		(0.0111)	(0.1785)	(2.6185)	(2.6412)	(0.1305)
	24 days	0.0103	0.2317	4.5532	4.6069	1.4550
		(0.0098)	(0.1740)	(1.0380)	(1.0754)	(0.1288)
$\hat{\sigma}_{aoc}$	5 days	−0.0388	−0.3285	23.2214	23.3293	3.2985
		(0.0105)	(0.2010)	(3.3425)	(3.2944)	(0.1312)
	24 days	0.0155	0.2772	4.8577	4.9345	1.5605
		(0.0107)	(0.2015)	(1.0953)	(1.1521)	(0.1358)
$\hat{\alpha}_p$	1 day	−0.1505	−1.6555	12.6511	15.3918	2.9704
		(0.0055)	(0.0718)	(0.8045)	(0.7431)	(0.0454)
	5 days	−0.1032	−1.1633	3.3992	4.7525	1.6505
		(0.0088)	(0.0622)	(0.5035)	(0.4931)	(0.0602)
	24 days	−0.0888	−1.0714	0.9521	2.1002	1.2088
		(0.0045)	(0.0785)	(0.1545)	(0.2655)	(0.0885)
$\hat{\sigma}_{GK}$	1 day	−0.1796	−2.0115	8.7125	12.7587	2.7085
		(0.0055)	(0.0450)	(0.4125)	(0.4952)	(0.0395)
	5 days	−0.1405	−1.6920	2.4813	5.3442	1.7997
		(0.0045)	(0.0611)	(0.2952)	(0.4315)	(0.0605)
	24 days	−0.1295	−1.6650	1.1405	3.9128	1.6751
		(0.0052)	(0.0885)	(0.2611)	(0.5352)	(0.0925)
$\hat{\sigma}_{RS}$	1 day	−0.1921	−2.3415	16.4885	21.9711	3.3052
		(0.0031)	(0.0522)	(0.4685)	(0.6658)	(0.0432)
	5 days	−0.1525	−1.8302	4.2545	7.6041	2.0285
		(0.0036)	(0.0695)	(0.3795)	(0.6602)	(0.0674)
	24 days	−0.1325	−1.7145	1.8525	4.7920	1.7810
		(0.0044)	(0.1202)	(0.3912)	(0.9052)	(0.1195)
$\hat{\sigma}_{ARS}$	1 day	−0.2545	−2.9831	16.1544	25.0534	3.5985
		(0.0048)	(0.0712)	(1.2188)	(1.4166)	(0.0582)
	5 days	−0.1988	−2.5145	4.6495	10.9723	2.6031
		(0.0049)	(0.0789)	(0.6188)	(0.8051)	(0.0714)
	24 days	−0.2003	−2.4540	2.2855	8.3077	2.5189
		(0.0062)	(0.1166)	(0.4855)	(0.8035)	(0.1102)
$\hat{\sigma}_{AGK}$	1 day	−0.2395	−2.7785	8.5755	16.2956	3.0995
		(0.0048)	(0.0645)	(1.2258)	(1.4964)	(0.0550)
	5 days	−0.2042	−2.5262	3.0556	9.4373	2.5789
		(0.0049)	(0.0788)	(0.6804)	(0.9911)	(0.0746)
	24 days	−0.1940	−2.4945	1.7045	7.9271	2.4822
		(0.0058)	(0.1322)	(0.5521)	(1.1049)	(0.1132)
$\hat{\sigma}_{YZ}$	5 days	−0.1228	−1.4922	2.6015	4.8282	1.7105
		(0.0043)	(0.0550)	(0.3113)	(0.4011)	(0.0498)
	24 days	−0.0997	−1.3748	1.1011	2.9912	1.3852
		(0.0047)	(0.0911)	(0.2633)	(0.4135)	(0.0666)

(*Continued*)

TABLE II (Continued)
Estimated Comparison Criteria

		Prop. Bias	*Bias*	*Variance*	*MSE*	*MAD*
		Panel B. Deutsche mark/Dollar exchange rate				
$\hat{\sigma}_{cc}$	1 day	−0.2589	−2.5922	40.7631	47.4826	5.5233
		(0.0095)	(0.1156)	(1.4978)	(1.3985)	(0.0722)
	5 days	−0.0970	−0.9808	10.0695	11.0315	2.6113
		(0.0104)	(0.1285)	(0.6722)	(0.6988)	(0.0821)
	24 days	−0.0605	−0.6052	2.2449	2.6112	1.3312
		(0.0111)	(0.1346)	(0.3015)	(0.3002)	(0.0786)
$\hat{\sigma}_{acc}$	5 days	−0.0998	−1.0436	12.6214	13.7105	2.8564
		(0.0095)	(0.1455)	(0.9180)	(0.9185)	(0.0923)
	24 days	−0.0602	−0.6001	2.5181	2.8783	1.3813
		(0.0089)	(0.1235)	(0.3245)	(0.3326)	(0.0825)
$\hat{\sigma}_{p}$	1 day	−0.1589	−1.5618	8.2012	10.6404	2.6281
		(0.0042)	(0.0435)	(0.3350)	(0.3104)	(0.0320)
	5 days	−0.1141	−1.2183	2.0021	3.4865	1.5367
		(0.0046)	(0.0546)	(0.1295)	(0.1712)	(0.0380)
	24 days	−0.0986	−1.1122	0.5102	1.7472	1.1642
		(0.0061)	(0.0581)	(0.0633)	(0.1489)	(0.0521)
$\hat{\sigma}_{GK}$	1 day	−0.1645	−1.6422	5.5206	8.2174	2.2988
		(0.0033)	(0.0380)	(0.2205)	(0.2497)	(0.0298)
	5 days	−0.1375	−1.3852	1.3918	3.3107	1.5222
		(0.0039)	(0.0462)	(0.1080)	(0.1749)	(0.0312)
	24 days	−0.1255	−1.3265	0.4108	2.1705	1.3189
		(0.0048)	(0.0512)	(0.0611)	(0.1680)	(0.0522)
$\hat{\sigma}_{RS}$	1 day	−0.1923	−1.9218	7.2435	10.9368	2.6147
		(0.0038)	(0.0377)	(0.2153)	(0.2629)	(0.0268)
	5 days	−0.1595	−1.6285	1.8223	4.4743	1.7922
		(0.0041)	(0.0451)	(0.1082)	(0.2192)	(0.0323)
	24 days	−0.1456	−1.5189	0.5089	2.8160	1.5237
		(0.0048)	(0.0521)	(0.0602)	(0.2289)	(0.0562)
$\hat{\sigma}_{ARS}$	1 day	−0.2257	−2.2618	7.0512	12.1669	2.7590
		(0.0041)	(0.0421)	(0.3812)	(0.4423)	(0.0341)
	5 days	−0.1947	−2.0011	1.8187	5.8231	2.0896
		(0.0049)	(0.0532)	(0.1587)	(0.2482)	(0.0432)
	24 days	−0.1791	−1.9256	0.5135	4.2214	1.9307
		(0.0058)	(0.0612)	(0.0632)	(0.2021)	(0.0588)
$\hat{\sigma}_{AGK}$	1 day	−0.2112	−2.0831	5.1112	9.4506	2.5103
		(0.0039)	(0.0439)	(0.3852)	(0.4785)	(0.0401)
	5 days	−0.1785	−1.8485	1.3085	4.7255	1.9152
		(0.0048)	(0.0521)	(0.1588)	(0.2912)	(0.0423)
	24 days	−0.1672	−1.7789	0.4098	3.5743	1.7785
		(0.0057)	(0.0632)	(0.0756)	(0.2784)	(0.0632)
$\hat{\sigma}_{YZ}$	5 days	−0.1445	−1.4601	1.4212	3.5531	1.6201
		(0.0046)	(0.0425)	(0.1088)	(0.1721)	(0.0388)
	24 days	−0.1305	−1.3618	0.3898	2.2443	1.3693
		(0.0053)	(0.0535)	(0.0485)	(0.1632)	(0.0541)

TABLE II (Continued)

		Prop. Bias	*Bias*	*Variance*	*MSE*	*MAD*
		Panel C. Japanese yen/Dollar exchange rate				
$\hat{\sigma}_{cc}$	1 day	−0.2851	−2.9909	50.1359	59.0814	5.9828
		(0.0085)	(0.1288)	(3.0577)	(2.9511)	(0.0828)
	5 days	−0.1431	−1.5198	13.1023	15.4121	3.0737
		(0.0098)	(0.1431)	(1.0652)	(0.9985)	(0.0937)
	24 days	−0.0955	−1.0032	2.9285	3.9349	1.6754
		(0.0117)	(0.1521)	(0.4156)	(0.3908)	(0.0932)
$\hat{\sigma}_{acc}$	5 days	−0.1412	−1.4635	15.6895	17.8313	3.3025
		(0.0105)	(0.1601)	(1.2741)	(1.1821)	(0.0925)
	24 days	−0.0951	−0.9921	3.1331	4.1174	1.6623
		(0.0121)	(0.1602)	(0.4875)	(0.4512)	(0.0987)
$\hat{\sigma}_{p}$	1 day	−0.1912	−2.0889	10.5576	14.9211	3.1222
		(0.0044)	(0.0502)	(0.7566)	(0.7025)	(0.0358)
	5 days	−0.1545	−1.6812	3.1498	5.9762	2.0422
		(0.0055)	(0.0618)	(0.3986)	(0.3612)	(0.0452)
	24 days	−0.1405	−1.5512	0.7225	3.1287	1.6125
		(0.0065)	(0.0732)	(0.1085)	(0.1895)	(0.0611)
$\hat{\sigma}_{GK}$	1 day	−0.2256	−2.2518	7.6322	12.7028	2.8901
		(0.0038)	(0.0488)	(0.5015)	(0.4811)	(0.0314)
	5 days	−0.1775	−1.8978	2.3402	5.9418	2.0912
		(0.0048)	(0.0556)	(0.4125)	(0.3741)	(0.0489)
	24 days	−0.1623	−1.7756	0.6185	3.7713	1.8285
		(0.0061)	(0.0685)	(0.1478)	(0.2085)	(0.0589)
$\hat{\sigma}_{RS}$	1 day	−0.2289	−2.4712	10.1201	16.2269	3.1735
		(0.0036)	(0.0412)	(0.4232)	(0.4458)	(0.0325)
	5 days	−0.1858	−2.0235	3.0115	7.1061	2.2561
		(0.0045)	(0.0547)	(0.3151)	(0.3312)	(0.0485)
	24 days	−0.1669	−1.8816	0.8256	4.3660	1.9312
		(0.0058)	(0.0558)	(0.0887)	(0.2621)	(0.0587)
$\hat{\sigma}_{ARS}$	1 day	−0.2662	−2.8426	9.6611	17.7415	3.3682
		(0.0044)	(0.0601)	(0.6502)	(0.6654)	(0.0421)
	5 days	−0.2212	−2.4685	2.7256	8.8191	2.6013
		(0.0052)	(0.0615)	(0.5154)	(0.4845)	(0.0356)
	24 days	−0.1995	−2.2585	0.7831	5.8839	2.2939
		(0.0066)	(0.0754)	(0.2132)	(0.2615)	(0.0612)
$\hat{\sigma}_{AGK}$	1 day	−0.2485	−2.6402	6.9889	13.9596	3.0898
		(0.0041)	(0.0521)	(0.5458)	(0.6458)	(0.0352)
	5 days	−0.2151	−2.3501	2.1002	7.6232	2.4618
		(0.0051)	(0.0621)	(0.4125)	(0.4498)	(0.0526)
	24 days	−0.1891	−2.1925	0.5502	5.3576	2.1829
		(0.0061)	(0.0701)	(0.1462)	(0.3185)	(0.0614)
$\hat{\sigma}_{YZ}$	5 days	−0.1817	−1.9766	2.4110	6.3190	2.1658
		(0.0049)	(0.0521)	(0.3889)	(0.3654)	(0.0418)
	24 days	−0.1512	−1.7418	0.6617	3.6956	1.7902
		(0.0062)	(0.0658)	(0.1587)	(0.2151)	(0.0621)

(*Continued*)

TABLE II (Continued)

Estimated Comparison Criteria

		Prop. Bias	Bias	Variance	MSE	MAD
		Panel D. British pound/Dollar exchange rate				
$\hat{\sigma}_{cc}$	1 day	−0.2732	−2.4212	38.5684	44.4306	5.3052
		(0.0095)	(0.1158)	(1.6360)	(1.4259)	(0.0454)
	5 days	−0.1125	−0.8998	11.3618	12.1714	2.8385
		(0.0112)	(0.1351)	(0.8189)	(0.7311)	(0.0725)
	24 days	−0.0605	−0.4228	3.4568	3.6356	1.6018
		(0.0125)	(0.1625)	(0.4158)	(0.3612)	(0.0878)
$\hat{\sigma}_{acc}$	5 days	−0.1315	−1.0212	13.3297	14.3725	3.0458
		(0.0121)	(0.1458)	(0.9825)	(0.8854)	(0.0815)
	24 days	−0.0601	−0.4195	3.6582	3.8342	1.6712
		(0.0167)	(0.1698)	(0.4221)	(0.3685)	(0.0825)
$\hat{\sigma}_{p}$	1 day	−0.1697	−1.4201	8.2211	10.2377	2.5802
		(0.0048)	(0.0458)	(0.3654)	(0.3085)	(0.0212)
	5 days	−0.1202	−1.0605	2.5785	3.7031	1.6459
		(0.0065)	(0.0631)	(0.1621)	(0.1658)	(0.0358)
	24 days	−0.0996	−0.8976	0.9376	1.7433	1.1249
		(0.0091)	(0.0815)	(0.1032)	(0.1625)	(0.0526)
$\hat{\sigma}_{GK}$	1 day	−0.1711	−1.4985	5.3256	7.5711	2.2645
		(0.0039)	(0.0365)	(0.1984)	(0.1967)	(0.0211)
	5 days	−0.1342	−1.2369	1.5335	3.0634	1.4902
		(0.0049)	(0.0477)	(0.0982)	(0.1265)	(0.0265)
	24 days	−0.1198	−1.1337	0.5708	1.8561	1.1881
		(0.0072)	(0.0623)	(0.0635)	(0.1658)	(0.0518)
$\hat{\sigma}_{RS}$	1 day	−0.1921	−1.7012	6.7701	9.6642	2.4712
		(0.0038)	(0.0376)	(0.1758)	(0.1955)	(0.0225)
	5 days	−0.1496	−1.3856	1.7002	3.6201	1.6141
		(0.0049)	(0.0412)	(0.0912)	(0.1654)	(0.0395)
	24 days	−0.1312	−1.2754	0.5502	2.1768	1.2959
		(0.0069)	(0.0632)	(0.0589)	(0.2018)	(0.0526)
$\hat{\sigma}_{ARS}$	1 day	−0.2189	−1.9989	6.4718	10.4674	2.5912
		(0.0042)	(0.0358)	(0.2845)	(0.3236)	(0.0258)
	5 days	−0.1905	−1.7996	1.5935	4.8320	1.9091
		(0.0052)	(0.0528)	(0.1214)	(0.1712)	(0.0395)
	24 days	−0.1708	−1.6585	0.5155	3.2661	1.6732
		(0.0068)	(0.0601)	(0.0615)	(0.1821)	(0.0578)
$\hat{\sigma}_{AGK}$	1 day	−0.2052	−1.8612	4.8402	8.3043	2.4105
		(0.0042)	(0.0452)	(0.2854)	(0.3331)	(0.0269)
	5 days	−0.1632	−1.5912	1.3708	3.9027	1.6989
		(0.0053)	(0.0332)	(0.1160)	(0.2110)	(0.0280)
	24 days	−0.1601	−1.5568	0.5075	2.9311	1.5527
		(0.0070)	(0.0631)	(0.0615)	(0.2254)	(0.0526)
$\hat{\sigma}_{YZ}$	5 days	−0.1398	−1.2631	1.5743	3.1697	1.4985
		(0.0052)	(0.0480)	(0.1132)	(0.1487)	(0.0312)
	24 days	−0.1192	−1.1288	0.5780	1.8522	1.1865
		(0.0074)	(0.0610)	(0.0652)	(0.1552)	(0.0512)

Note. The table presents the proportional bias, bias, variance, mean squared error (MSE), and mean absolute deviation (MAD) of various volatility estimators when volatility is estimated with the use of daily data over periods of 1, 5, and 21 (trading) days. Numbers in parentheses are bootstrapped standard errors. The traditional volatility estimator assuming no drift is denoted $\hat{\sigma}_{cc}$; $\hat{\sigma}_{acc}$ adjusts for a possible drift; $\hat{\sigma}_{p}$ denotes the Parkinson estimator; $\hat{\sigma}_{GK}$ denotes the Garman and Klass estimator; $\hat{\sigma}_{RS}$ denotes the Rogers and Satchell estimator. $\hat{\sigma}_{ARS}$ and $\hat{\sigma}_{AGK}$ are versions of, respectively, the Garman and Klass and Rogers and Satchell estimators that adjust for discrete trading, and $\hat{\sigma}_{YZ}$ denotes the Yang and Zhang estimator. The sample periods are January 3, 1989–December 30, 2003 for the S&P 500 futures contract, and December 3, 1986–August 30, 2003 for the exchange rates.

volatility estimators are almost always significantly less biased than the traditional estimator. For example, for the Deutsche mark/dollar exchange rate, the bias of the Garman and Klass estimator is less than two thirds that of the traditional estimator. The only exception to this is the bias associated with the Rogers and Satchell adjusted extreme-value estimators, $\hat{\sigma}_{\mathrm{ARS}}$ and $\hat{\sigma}_{\mathrm{AGK}}$. In all cases, these adjustments make extreme-value estimators more biased, not less, and in the case of the S&P futures contract, the increase in bias is sufficient to make $\hat{\sigma}_{\mathrm{ARS}}$ and $\hat{\sigma}_{\mathrm{AGK}}$ more biased than the traditional estimator.

Second, extreme-value estimators provide substantial efficiency gains. Interestingly, these efficiency gains have empirical magnitudes that are very close to their theoretical values. Taking again the example of the Deutsche mark/dollar exchange rate, the estimated efficiency gain of the Garman and Klass estimator is 7.4 (=40.7631/5.5206), which is exactly its theoretical value.

Third, across all four data sets, all extreme-value volatility estimators outperformed the traditional estimator. This result obtains for both comparison criteria (mean-squared error and mean absolute deviation).

Fourth, of practical interest, a single extreme-value estimator, that of Garman and Klass, outperforms the others across the four data sets.

Weekly and Monthly Volatility Estimation

At the weekly and monthly frequencies, it was found that extreme-value volatility estimators are more biased than the traditional estimator. As the length of the estimation window increases (from a day, to a week, to a month), the bias of the traditional estimator decreases much more than that of the extreme-value estimators. As in the case of daily volatility estimation, extreme-value estimators were found to yield large efficiency gains; again, these efficiency gains have estimated values that are very close to their theoretical counterparts.

At the weekly frequency, it is still the case that all extreme-value volatility estimators outperform the traditional estimator, but this does not hold in the case of monthly volatility estimation: at the monthly frequency, the traditional estimator is significantly less biased than the extreme-value estimators, and it is difficult for extreme-value estimators to make up for this high bias with smaller variance. Still, the best-performing extreme-value estimators, namely, those of Garman and Klass, Parkinson, and Yang and Zhang, do outperform the traditional estimator, at both the weekly and monthly frequencies, in both a mean-squared-error and a mean absolute deviation sense.

CONCLUSION

This article compares the bias and efficiency of a number of extreme-value volatility estimators that have appeared over the past two decades. To get around the difficulty that volatility is not directly observable, recent advances in econometric methodology and newly available data are used. Specifically, high-frequency data on very actively traded assets (S&P 500 futures and three currencies) are used to construct measures of realized volatility, and then the performances of the various volatility estimators are compared against this benchmark.

Several important results emerge from this analysis. Strong support is found for the use of extreme-value volatility estimators when estimating daily volatilities: At the daily frequency, extreme-value estimators are less biased than the traditional estimator and they are also significantly more efficient. Of practical interest, the estimator of Garman and Klass is found to be the single best-performing estimator, but in fact all extreme-value estimators outperform the traditional estimator in a mean-squared-error and mean absolute deviation sense.

At the weekly and monthly frequencies, the results also indicate that extreme-value estimators are appropriate, but the evidence is more mixed. This is because extreme-value estimators are more biased than the traditional estimator at those frequencies, so the efficiency gains, while large, may not be sufficient to offset the large biases. To gain intuition about the sharp decrease in the bias of the traditional estimator as the length of the estimation window increases (from a day, to a week, to a month), recognize that in order to estimate volatility over a given month, the traditional estimator does not merely square the monthly return (which, in effect, is what it does at the daily frequency), but rather divides the 1-month interval into 21 periods and sums those squared daily returns (which goes in the direction of the continuous-record asymptotics that realized volatility relies upon). From a practical perspective, although the traditional estimator performs relatively well at the weekly and monthly frequencies, there are three extreme-value volatility estimators (those of Garman and Klass, Parkinson, and Yang and Zhang) that consistently perform better.

BIBLIOGRAPHY

Andersen, T. G., & Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. International Economic Review, 39, 885–905.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Ebens, H. (2001). The distribution of realized stock return volatility. Journal of Financial Economics, 61, 43–76.
Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2001). The distribution of exchange rate volatility. Journal of the American Statistical Association, 96, 42–55.
Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (in press). Modeling and forecasting realized volatility. Econometrica.
Andreou, E., & Ghysels, E. (2002). Rolling-sample volatility estimators: Some new theoretical, simulation, and empirical results. Journal of Business and Economic Statistics, 20, 363–376.
Bali, T. G. (2003). An extreme value approach to estimating volatility and value at risk. Journal of Business, 76, 83–108.
Bali, T. G., & Neftci, S. (2003). Disturbing extremal behavior of spot rate dynamics. Journal of Empirical Finance, 10, 455–477.
Ball, C. A., & Torous, W. N. (1984). The maximum likelihood estimation of security price volatility: Theory, evidence, and application to option pricing. Journal of Business, 57, 97–112.
Barndorff-Nielsen, O. E., & Shephard, N. (2002). Econometric analysis of realised volatility and its use in estimating stochastic volatility models. Journal of the Royal Statistical Society, Series B, 64, 253–280.
Beckers, S. (1983). Variances of security price returns based on high, low, and closing prices. Journal of Business, 56, 97–112.
Brown, G. W., & Hartzell, J. C. (2001). Market reaction to public information: The atypical case of the Boston Celtics. Journal of Financial Economics, 60, 333–370.
Cho, D. C., & Frees, E. W. (1988). Estimating the volatility of discrete stock prices. Journal of Finance, 43, 451–466.
Efron, B., & Tibshirani, R. (1986). Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. Statistical Science, 1, 54–75.
Engle, R. F., & Patton, A. (in press). Impacts of trades in an error-correction model of quote prices. Journal of Financial Markets.
Garman, M. B., & Klass, M. J. (1980). On the estimation of security price volatilities from historical data. Journal of Business, 53, 67–78.
Kunitomo, N. (1992). Improving the Parkinson method of estimating security price volatilities. Journal of Business, 65, 295–302.
Lo, A. W., & Wang, J. (1995). Implementing option pricing models when asset returns are predictable. Journal of Finance, 50, 87–129.
Marsh, T., & Rosenfeld, E. (1986). Non-trading, market making, and estimates of stock price volatility. Journal of Financial Economics, 15, 359–372.
Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. Journal of Financial Economics, 3, 125–144.
Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. Journal of Business, 53, 61–65.
Rogers, L. C. G., & Satchell, S. E. (1991). Estimating variance from high, low and closing prices. Annals of Applied Probability, 1, 504–512.

Rogers, L. C. G., Satchell, S. E., & Yoon, Y. (1994). Estimating the volatility of stock prices: A comparison of methods that use high and low prices. Applied Financial Economics, 4, 241–247.

Wiggins, J. B. (1991). Empirical tests of the bias and efficiency of the extreme-value variance estimator for common stocks. Journal of Business, 64, 417–432.

Yang, D., & Zhang, Q. (2000). Drift independent volatility estimation based on high, low, open, and close prices. Journal of Business, 73, 477–492.